Claims

1. A method performed on a programmable digital computer for processing relational databases having data files, comprising:

      defining a schema of the databases as a script;

      compiling the script into a representation of the schema;

      constructing an executable application program for processing the databases;

      installing the representation of the schema in the executable application program such that the schema is stored separately from the data file of the relational databases processed by the application program.

2. A medium containing instructions and data for executing the method of claim 1 on a programmable digital computer.

3. A method according to claim 1, wherein the schema includes definitions of tables, columns for the tables, and data types for the columns.

4. A method according to claim 1, wherein the schema further includes at least one index.

5. A method according to claim 3, further comprising compiling the schema script into a helper file containing structural definitions for the tables and a set of macros.

6. A method according to claim 1, further comprising compiling the script into a further standalone representation of the schema independent of the first representation.

7. A method according to claim 6 further comprising installing the standalone representation in a user computer independently of the application program.

8. A method performed on a programmable digital computer useful in processing relational databases with an application program, comprising:

      recording a desired schema for the databases as a human-readable script in a source format;

compiling the script into a representation capable of being included in the code of the application program, such that the schema forms a part of the application program rather than a part of the databases.

9. A medium containing instructions and data for executing the method of claim 8 on a programmable digital computer.

10. A method according to claim 8, wherein the representation includes a globally unique schema identifier.

11. A method according to claim 10, wherein the representation further includes a version number.

12. A method according to claim 8, wherein the schema includes definitions of tables and columns for the tables.

13. A method according to claim 12, wherein the schema further includes at least one index.

14. A method according to claim 12, wherein the representation includes structural constructs representing the definitions of the tables and columns.

15. A method performed on a programmable digital computer for processing relational databases with an application program, comprising:

    selecting one of the databases for processing;

    accessing a data file of the selected database;

    reading a pointer in the data file indicating a schema file stored independently of the data file;

    accessing the schema file separately from the data file;

    accessing database data within the data file in accordance with the schema file.

16. A medium containing instructions and data for executing the method of claim 15 on a programmable digital computer.

17. A method according to claim 15, wherein the pointer contains an identifier for locating the schema file.

18. A method according to claim 17, wherein the pointer further contains a version identifier.

19. A method performed on a programmable digital computer for processing relational databases having data files organized according to a defined schema, comprising:

defining application program code for performing at least one function upon the databases;

including with the application code a schema file separate from the data files and representing the schema of the data files, such that the function operates upon the data files in accordance with the schema.

20. A medium containing instructions and data for executing the method of claim 19 on a programmable digital computer.

21. A method according to claim 19, wherein the application code accesses a further program for executing the function, and wherein the application code passes the schema file to the further code for locating particular data in the data files.

22. A method according to claim 21, wherein the further function is a database engine.

23. A method according to claim 21, wherein the schema file includes structural constructs representing definitions of the tables and columns in the schema.

24. A method according to claim 1, wherein the schema file further includes at least one index.

25. A method according to claim 23, wherein the constructs are separately integrated into the application code.

26. A system for processing a relational database, comprising:
        a data file for containing data for the database according to a schema;
        a schema definition for the database;

a compiler responsive to the schema definition for producing a representation of the schema independent from the data file.

27. A system according to claim 6, wherein the schema definition is a human-readable file.

28. A system according to claim 26, wherein the representation is a binary representation.

29. A system according to claim 26, wherein the schema definition includes specifications of tables and columns.

30. A system according to claim 29, wherein the representation includes constructs embodying the specifications of the tables and columns.

31. A system of relational databases, comprising:

a plurality of schema files each representing the schemata of different ones of the databases;

a plurality of data files separate from the schema files and each containing relational data organized according to a corresponding one of the schemata;

a pointer in each of the data files identifying one of the schema files representing the corresponding schema.

32. A system according to claim 31, further comprising an application program for selectively accessing any one of the data files and for accessing that one of the schema files identified by the pointer in the accessed data file.

33. A system according to claim 32, further comprising a database engine coupled to the application program for performing an operation on the one data file in accordance with the one schema file accessed by the application program.

34. A system for processing a relational database, comprising:

a data file containing relational data organized according to a schema;

an application program separate from the data file and including

a representation of the schema of the data file,

code responsive to the representation for performing an operation on the data file in accordance with the schema.

35. A system according to claim 34, further comprising a database engine coupled to the application program for executing the operation upon the data file in response to receiving the representation of the schema from the application program.

36. A system for processing a relational database, comprising:

a data file containing relational data organized according to a schema;

a schema file separate from the data file and containing constructs representing the schema;

an application program separate from the data file and including at least one of the constructs taken from the schema representation and further including code for manipulating data according to the at least one construct.

37. A system according to claim 36, wherein the schema file further includes code for manipulating the at least one construct taken from the schema.

38. A system according to claim 37, wherein at least some of the code in the application program is taken from the schema file.

39. A relational database, comprising:

a data file for containing relational data for the database according to a schema;

a schema file entirely separate from the data file for representing the schema of the relational data to an application program that accesses the data file.